

Rapid decisions from experience



Matthew D. Zeigenfuse^{a,*}, Timothy J. Pleskac^b, Taosheng Liu^b

^a Psychologisches Institut, Universität Zürich, Binzmühlestrasse 14 Box 26, CH-8050 Zürich, Switzerland

^b Department of Psychology, Michigan State University, 316 Physics Road, East Lansing, MI 48864, United States

ARTICLE INFO

Article history:

Received 2 March 2013

Revised 19 December 2013

Accepted 20 December 2013

Keywords:

Decision making

Decisions from experience

Rank-dependent utility theory

Prospect theory

Sequential sampling

ABSTRACT

In many everyday decisions, people quickly integrate noisy samples of information to form a preference among alternatives that offer uncertain rewards. Here, we investigated this decision process using the Flash Gambling Task (FGT), in which participants made a series of choices between a certain payoff and an uncertain alternative that produced a normal distribution of payoffs. For each choice, participants experienced the distribution of payoffs via rapid samples updated every 50 ms. We show that people can make these rapid decisions from experience and that the decision process is consistent with a sequential sampling process. Results also reveal a dissociation between these preferential decisions and equivalent perceptual decisions where participants had to determine which alternatives contained more dots on average. To account for this dissociation, we developed a sequential sampling rank-dependent utility model, which showed that participants in the FGT attended more to larger potential payoffs than participants in the perceptual task despite being given equivalent information. We discuss the implications of these findings in terms of computational models of preferential choice and a more complete understanding of experience-based decision making.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Imagine you are a day trader deciding whether or not to buy a stock. To make a choice, you watch market data moving along LED ribbons around trading floors or at the bottom of computer monitors trying to gauge the price trend. The information flows by in real time so to decide if the market is trending up or down you have to quickly integrate the information racing by. How does the day trader make this preferential choice where there is not an objectively correct answer, but instead where the choice is based on some subjective value of the alternatives? One reasonable hypothesis is that the day trader makes a decision using a sequential sampling process (Busemeyer & Townsend, 1993; Usher & McClelland, 2004). During this process, as decision makers deliberate, they sequentially

sample payoff information about the possible alternatives – either directly from the alternatives or from their memory of past experience with the alternatives – and integrate that information over time. Once a threshold of evidence is reached a decision is made accordingly.

Despite the plausibility of this sequential sampling hypothesis, we do not know how well decision makers integrate rapid samples of payoff information, which is a critical assumption of the process. More generally, it is an open question how well a sequential sampling process describes choices when decision makers have to rapidly process payoff information as in the case of our day trader. Most of the empirical work to date on risky decision making has focused on how people choose between monetary lotteries that are presented as static, symbolic descriptions of the payoffs and probabilities (i.e., decisions from description) (Weber, Shafir, & Blais, 2004). Other studies have used the so-called decisions-from-experience paradigm where participants explicitly sample from the alternatives and

* Corresponding author. Tel.: +41 0446357374.

E-mail address: m.zeigenfuse@psychologie.uzh.ch (M.D. Zeigenfuse).

receive feedback over an extended number of trials (Hertwig & Erev, 2009). However, both of these decisions are quite different from the choices people make when they must integrate rapidly arriving payoff information.

Understanding how people accumulate payoff information in these situations may also improve our understanding of how people make decisions from description. Some computational models of these decisions assume a latent sequential sampling process drives the choice (e.g., Busemeyer & Townsend, 1993). In these models, when people are presented with a choice between a sure thing and a risky gamble it is assumed that they mentally simulate random samples from the gamble. These mentally simulated samples are then sequentially accumulated to a threshold that in turn determines the choice and response time. These models not only give a good account of overall choice patterns (Busemeyer & Townsend, 1993), but they also appear to describe the dynamics of deliberation as attention switches back and forth between attributes and alternatives (Busemeyer, 1985; Diederich & Busemeyer, 1999; Krajbich, Armel, & Rangel, 2010; Krajbich, Lu, Camerer, & Rangel, 2012; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010). Again, however, a critical question remains as to how and how well decision makers can integrate rapid samples of payoff information.

Indeed work in perceptual decision making suggests that humans and other primates can rapidly integrate changing perceptual information to make a decision and that this process is consistent with a sequential sampling process (Gold & Shadlen, 2007; Palmer, Huk, & Shadlen, 2005; Ratcliff & McKoon, 2008). Thus, not only is the assumption for rapid integration of payoff information plausible, but it seems a similar process is used to make perceptual decisions as has been posited for preferential decisions. This parallelism has, in fact, led to the intriguing hypothesis that preferential and perceptual choice may even use the same or similar cognitive machinery (Busemeyer, Jessup, Johnson, & Townsend, 2006; Shadlen, Kiani, Hanks, & Churchland, 2008; Summerfield & Tsetsos, 2012; Symmonds & Dolan, 2012).

To examine whether a sequential sampling process underlies rapid decisions from experience, and to test the idea that perceptual and preferential decisions are based on similar cognitive mechanisms, we developed a novel preferential decision task, which we call the Flash Gambling Task (FGT). The FGT uses dynamic dot stimuli like those in perceptual decision tasks. In the FGT, participants make a series of decisions between a certain and an uncertain alternative, with the payoff amount indicated by the number of dots in the display. Importantly, outcomes from the uncertain alternative are dynamically updated every 50 ms via draws from an unknown payoff distribution. We compared the FGT to its perceptual analog, where participants were shown the same stimuli and told to identify which option had the higher average number of dots. This comparison allowed us to directly test the hypothesis that people use the same or similar process to make preferential decisions as they use to make perceptual decisions. As we will show, we found both similarities and important differences in the performance on the two tasks. To account for these results, we develop a computational model

that incorporates sequential sampling and rank dependent utility theory (Luce, 2000; Quiggin, 1982; Tversky & Kahneman, 1992).

2. Methods

2.1. Participants

Twenty-three students from Michigan State University completed the Flash Gambling Task (FGT) and twenty-five students completed the matched perceptual task. Participants in both tasks were paid \$8.00 for participation, which took one hour. They were also paid a bonus based on their choices. In the FGT, this was the reward earned on four randomly selected trials. In the perceptual task, it was \$0.005 times the number of times their response was correct. Both tasks were calibrated to have an average bonus of \$2.00.

2.2. Flash Gambling task

Stimuli were generated in MATLAB using MGL (<http://justingardner.net/mgl>) and displayed on a LCD monitor. Two circular displays of white dots were shown on a black background. Each display had a diameter of 12° of visual angle. One was located 9° to the left of a green central fixation and the other 9° to the right. The location of the certain alternative was fixed for a given participant but counterbalanced across participants.

The certain alternative was filled with randomly placed dots that remained throughout the duration of a trial. The uncertain alternative had a dynamic display of dots that changed every 50 ms (20 Hz). At each update, the number of dots was drawn from a normal distribution, and then shown in random locations in the display. We truncated the normal distribution at ± 2 standard deviations to ensure the uncertain alternative always contained at least 30 dots. Participants were told to choose the option they would like to receive a draw from. After making a choice,




Fig. 1. Illustration of the stimuli in the flash gambling task.

participants received the payoff from the next draw, which was displayed as numerical feedback (Fig. 1).

2.3. Perceptual task

The perceptual task was identical to the FGT except participants were told to identify the option that they believed showed the higher average number of dots. Participants received feedback if they were correct or not.

2.4. Design & procedure

Participants in both tasks completed ten blocks of 84 trials. In addition to the between-subjects perceptual versus gambling manipulation, we manipulated three factors within subjects: the mean difference in the number of dots between the certain and uncertain alternative, the standard deviation in the number of dots in the uncertain alternative, and the number of dots in the certain alternative. Each of these manipulations occurred between trials within a block.

The manipulation of the mean difference in the number of dots (uncertain offset) allowed us to determine whether participants could discriminate differences between the certain and uncertain alternatives. We used three levels of uncertain offset, with the uncertain option having 30 fewer, an equal number, and 30 more dots than the certain alternative. We manipulated the standard deviation of the uncertain alternative to investigate a possible role of payoff variance on choice (Bussemeyer, 1985). The standard deviation in the uncertain alternative was either 15 or 50 dots. Finally, we varied the number of dots in the certain alternative to encourage participants to attend to both the certain and uncertain alternative and treat them as separate options. The certain alternative contained either 130 or 250 dots.

During a trial, participants were shown the central fixation first for 200 ms, after which the certain and uncertain alternatives appeared and remained on screen until response. Participants chose the preferred option by pressing one of two keys on a keyboard. Once a response was made, the fixation dot was replaced with feedback. Feedback in the FGT was the value of a random draw from the chosen alternative. Feedback in the perceptual task was whether participant's decision was correct.

3. Results

3.1. Behavioral results

We analyzed data from twenty FGT and twenty-four perceptual participants. Three FGT participants were excluded for having over 10% (84 trials) of their response times less than 100 ms, which is faster than the typical time it takes to simply press a button. One perceptual participant was excluded for performing the task incorrectly. In all behavioral analyses, we report main effects and interactions between all the manipulated variables. Unless specified otherwise, we used one-sided *t*-tests and adjusted the degrees of freedom to maintain a family wise error rate

of $\alpha = 0.05$ for all post hoc contrasts we investigated for a given analysis. These contrasts were performed using the `multcomp` package in R (Hothorn, Bretz, & Westfall, 2008).

As previously mentioned, the number of dots in the certain alternative was varied in order to encourage participants to treat the certain alternative as a separate gamble. We did not, however, scale the magnitude of the uncertain offset or standard deviation. As a result, consistent with Weber's law, participants in both tasks were not able to discriminate differences between the two alternatives in the 250 dot condition. Consequently, though we report main effects and interactions involving the number of certain dots, in figures, post hoc analyses, and subsequent computational model-based analyses, we will only use trials where the certain alternative contained 130 dots.

3.1.1. Choice

Panels A and B of Fig. 2 show the proportion of trials participants chose the uncertain alternative as a function of uncertain offset and standard deviation and task. Because our choice data was binary, we employed a logistic mixed model to determine the effects of the between-subjects factor of task (gambling or perceptual) and the within-subjects factors of uncertain offset, standard deviation, and the number of certain dots.

The logistic model revealed a main effect of uncertain offset ($\chi^2(2) = 64.5, p < 0.001$). Specifically, the proportion of trials participants chose the uncertain alternative increased as the mean difference went from favoring the certain to the uncertain alternative, suggesting participants were able to discriminate which alternative had more dots. We confirmed this by testing whether the difference in the log-odds of choosing the uncertain alternative between 0 and -30 uncertain offset trials was positive ($z = 14.8, p_{adj} < 0.001$) and whether the corresponding difference between $+30$ and 0 uncertain offset trials was positive ($z = 12.8, p_{adj} < 0.001$).

The logistic model also revealed a main effect of task, with participants choosing the uncertain alternative more often in the FGT than in the perceptual analog ($\chi^2(1) = 13.4, p < 0.001$), as well as an interaction between task and uncertain offset ($\chi^2(1) = 16.4, p < 0.001$). Panels A and B of Fig. 2 show that this interaction arose because the uncertain alternative choice proportions changed more in the perceptual task than in the FGT, suggesting differences in the decision process between the two tasks. We confirmed that the log-odds difference between 0 and -30 uncertain offset trials was larger in the perceptual task than in the FGT ($z = 4.08, p_{adj} < 0.001$), but that the analogous difference between $+30$ and 0 trials was not ($z = 1.69, p_{adj} = 0.395$). The former difference appears to arise from the uncertain alternative being chosen more often in the FGT than the perceptual task for -30 uncertain offset trials, since this log-odds difference was positive ($z = 3.6, p_{adj} = 0.001$), but the corresponding difference in the $+30$ offset condition was not ($z = 1.5, p_{adj} = 0.475$). In other words, this interaction was due in part to a disproportionate preference for the uncertain alternative in the -30 offset condition in the gambling task. Our computational modeling will examine this disproportionate preference in more depth.




Fig. 2. (A) and (B) Choice proportion for the uncertain alternative on the 130 certain dot trials. Panel A: standard deviation of 15 dots in the uncertain alternative, Panel B: standard deviation of 50 dots. Better alternative is defined by the difference in mean number of dots between certain and uncertain alternative (+30: uncertain > certain, 0: uncertain = certain, -30: uncertain < certain). (C) and (D) Mean response time on 130 certain dot trials. Panel C: standard deviation of 15 dots in the uncertain alternative, Panel D: standard deviation of 50 dots. Gray bars: FGT data, white bars: perceptual data. Error bars correspond to one standard error of the mean (SEM).

We also found a two-way interaction between the task and the number of certain dots ($\chi^2(1) = 26.5, p < 0.001$) and a three-way interaction between task, uncertain offset and the number of dots ($\chi^2(2) = 12.8, p = 0.001$), though these are not shown in Fig. 2. As discussed earlier, these interactions were the result of participants not being able to discriminate between the uncertain and certain alternative in the 250 dots condition.

3.1.2. Response time

Panels C and D of Fig. 2 show the response time data. We examined differences in RTs using a repeated-measures ANOVA of $\ln(RT)$ on the factors task, uncertain offset, standard deviation and the number of certain dots. We used $\ln(RT)$ rather than RT to compensate for the skewed distribution of RTs. Due to the unbalanced design, we used a multi-level model and report likelihood ratio tests.

The ANOVA reveals a main effect of uncertain offset ($\chi^2(2) = 21.8, p < 0.001$). There was also a main effect of task ($\chi^2(1) = 10.3, p = 0.001$), with faster RTs in the FGT than in the perceptual analog. Finally, we found a main effect of the number of certain dots ($\chi^2(1) = 22.4, p < 0.001$) as well as an interaction between task and the number of certain dots ($\chi^2(2) = 11.1, p < 0.001$).

3.2. Computational modeling results

To better compare people's decision processes in the FGT and perceptual task, we fit a drift diffusion model to

the data from both tasks, which models the decision process as a sequential sampling process (e.g., Laming, 1968; Link & Heath, 1975; Ratcliff, 1978; Stone, 1960). The drift diffusion model is a mathematical formulation of a sequential sampling process where participants are assumed to sequentially sample noisy information and accumulate it as evidence until a threshold is reached initiating a response. As a mathematical model we can fit it to the observed choices and response time distributions to decompose observed behavior into four psychologically meaningful parameters: (a) *drift rate* indexing the average amount and direction of evidence accumulated; (b) *decision threshold* indexing the amount of evidence required to make a decision; (c) *bias* indexing the prior bias in favor of choosing the uncertain or certain alternative; and (d) *non-decision time* indexing the amount of time spent on decision-irrelevant processing.

The model does not use the observed stimulus information to drive the evidence accumulation process; rather, it treats the evidence accumulation process as latent. The drift rate, in turn, estimates the rate and direction at which evidence would be accumulated in order to account for the observed choice data. Thus, we were particularly interested in whether the difference in choice proportions and RTs between the FGT and perceptual analog would also result in differences in drift rate. Such a drift difference would indicate that the gamble frame changes the evidence that subjects accumulated. Differences between the two tasks could also arise in the other parameters.

For instance, one might speculate that the difference in RTs between the two tasks could be due to a greater level of response caution during the perceptual task. If this were true, then greater response caution should manifest itself via higher decision thresholds allowing participants to collect more evidence for each decision.

We fit the drift diffusion model at the individual level using quantile maximum probability method (Heathcote, Brown, & Mewhort, 2002). This method fits the model to the quantiles of the observed response time distributions for uncertain and certain alternative choices (see also Busemeyer & Diederich, 2010). For each individual, regardless of whether they were in the FGT or perceptual condition, we fit a drift rate for each level of uncertain offset (collapsing across the different levels of standard deviation). Threshold, bias and non-decision time were set constant across offset and standard deviation conditions. As shown by Fig. 3, the model does a good job in recreating the choice proportion and RT data. The average mean-squared prediction error (MSE) over participants of the predicted choice proportions and expected RTs supports the conclusion of a good fit (average choice MSE = 0.028, average RT MSE = 0.112 s).

3.2.1. Drift rate

Fig. 4A plots average drift rates as a function of task and uncertain offset. We fit a linear mixed model with one between-subjects factor for task and one within-subjects factor for uncertain offset. We found a significant main effect of uncertain offset ($\chi^2(2) = 44.9, p < 0.001$), indicating that changes in the uncertain offset changed the rate of evidence accumulation. This difference in drift rates accounts for the effect of uncertain offset on choice proportions and RTs.

More importantly for this study is the difference in drift rate between the FGT and the perceptual analog, manifested as a main effect of task ($\chi^2(1) = 8.89, p = 0.002$), with participants accumulating evidence favoring the uncertain alternative faster in the FGT than in the perceptual task ($z = 2.77, p_{adj} = 0.007$). Moreover, we found that drift rates were positive on -30 offset trials in the FGT ($z = 2.13, p_{adj} = 0.041$), but were negative in the same trials in the perceptual task ($z = 2.09, p_{adj} = 0.045$). That is, when participants viewed the equivalent statistical

information but framed as payoffs, information processing changed in a way that led participants to prefer the uncertain alternative even when the objective evidence supported the certain alternative offering a higher expected reward. Finally, we also found a significant interaction between task and uncertain offset in the drift rates ($\chi^2(2) = 6.9, p = 0.031$). Since the interaction is scale-dependent, we refrain from further interpretation (Loftus, 1978; Wagenmaker, Krypotos, Criss, & Iverson, 2012).

3.2.2. Bias, threshold & non-decision time

We also examined whether participants in the FGT and perceptual task varied in the amount of evidence they required to make a decision, their response bias toward either alternative and the amount of time they spent on non-decisional processing (Fig. 4B–D). We tested each hypothesis using a two-sided *t*-test. In all cases, the results were consistent with chance variation (for threshold, $t(23.4) = -0.529$; for bias, $t(20.9) = 1.26$; for non-decision time, $t(34.7) = 0.852$, all *ps* > 0.2).

These results indicate that the observed differences in choice and RT across the FGT and perceptual task derive primarily from differences in the evidence being accumulated, and not from, for example, an overall bias for participants to be risk-seeking or for participants to set a larger threshold in the perceptual task. The differences in drift rates suggest that on average participants in the FGT collected evidence favoring the uncertain alternative faster than participants in the perceptual task, which in turn led to FGT participants responding faster and preferring the uncertain alternative more often. Moreover, the interaction between task and uncertain offset in the drift rate implies the critical difference between the FGT and the perceptual task is a difference in the evidence that is accumulated. In the discussion, we argue that this difference in drift rates between the FGT and its perceptual analog is due to participants' valuation process during the FGT. In particular, participants appear to weigh experienced events in an optimistic manner.

4. Discussion

In this paper, we investigated how people make rapid decisions from experience when they are required to




Fig. 3. Drift diffusion model fit and data. (A) Observed choice proportion for the uncertain alternative (bars), averaged across 15 and 50 dots standard deviation conditions. The model fit was shown as dots. (B) Observed mean response time (bars) averaged across standard deviation conditions and predicted mean response time (dots). Gray bars: FGT data, white bars: perceptual data. Error bars correspond to one standard error of the mean (SEM).




Fig. 4. Average parameter values of the diffusion model fits. (A) Mean drift rate. (B) Mean bias. (C) Mean threshold. (D) Mean non-decision time. Error bars are correspond to a single SEM.

quickly integrate a stream of information to make a choice. We showed that as rewards from an uncertain alternative became, on average, more attractive, preference shifted towards the uncertain alternative. Consistent with sequential sampling accounts of preferential choice (Bussey & Townsend, 1993; Usher & McClelland, 2004), choice and response times were well described by a drift diffusion process. This was also true in the perceptual analog of the task, suggesting these two different decisions use the same or similar systems (Bussey et al., 2006; Shadlen et al., 2008; Summerfield & Tsetsos, 2012; Symmonds & Dolan, 2012).

However, the study also revealed a dissociation between the two types of decisions. In particular, preference for the uncertain alternative was stronger and response times were faster during the FGT than in the perceptual analog. Our drift diffusion analyses suggest these results are not due to a response bias or differences in the location of the choice threshold, but the result of a difference in the evidence that is accumulated during deliberation. In other words, preferential choice under uncertainty appears to lead to a different representation of statistical information as compared to when participants made perceptual judgments under uncertainty. Given that participants in both conditions saw the exact same statistical information, the next question is what explains this difference in the accumulated evidence.

Behavioral decision theory offers several possible explanations for how the subjective representation of value can

depart from the objective value. In the remainder of this section, we determine whether three of the most common of these explanations – a utility function (Bernoulli, 1954), serial order effects (Hogarth & Einhorn, 1992) and rank dependent weights from rank-dependent utility theory (Luce, 2000; Quiggin, 1982; Tversky & Kahneman, 1992) – could account for our data. To do so, we show how each concept from behavioral decision theory morphs the evidence that is accumulated in a sequential sampling process and thus how it affects the drift rate.

4.1. The sequential sampling process model

Before developing the possible explanations, it is necessary to formally define the sequential sampling framework, which our analyses suggest provides a good description of the basic decision process. In both conditions, if decision makers were accumulating evidence objectively comparing the value of the certain alternative c to each of the k values sampled from the uncertain alternative, labeled $(y_1, \dots, y_k)$, then the accumulated evidence would be

$$E_k^{\text{obj}} = \sum_{j=1}^k (y_j - c). \quad (1)$$

Eq. (1) formally links the information presented to the participant to the evidence that is accumulated during the sequential sampling process, making the latent process

assumed in our earlier drift–diffusion model analyses explicit. Previous work (Palmer et al., 2005; Ratcliff & McKoon, 2008), as well as data from our perceptual task, suggest this is a reasonable approximation of the accumulated evidence during our perceptual condition. Unfortunately, the sequential sampling process in Eq. (1) cannot describe the decision process in the FGT because it cannot predict a positive drift rate when the uncertain offset is negative (see Fig. 4 Panel A). However, given its ability to account for perceptual data, we will build on Eq. (1) to incorporate effects from behavioral decision theory.

4.2. Sequential sampling of subjective expected utility

The first possible explanation we investigate is a utility function. Utility functions are a standard way of modeling the representational change between objective monetary amounts and subjective values. This function formalizes the idea that decision makers evaluate monetary outcomes based on their subjective value or utility rather than their objective value (Bernoulli, 1954; Kahneman & Tversky, 1979; Savage, 1954; von Neumann & Morgenstern, 1947). In sequential sampling models, utility impacts what evidence is accumulated and thus affects the measured drift rate (Busemeyer & Townsend, 1993). Decision makers typically exhibit decreasing sensitivity to payoffs as the magnitude of payoffs increase, which is modeled with a power function $u(x) = x^\theta$ with $0 < \theta < 1$ (Kahneman & Tversky, 1979; Luce, 2000). Substituting such a function into the evidence accumulation process (Eq. (1)) results in

$$E_k^{\text{util}} = \sum_{j=1}^k (y_j^\theta - c^\theta). \quad (2)$$

For utility to account for preferences, there must be a $\theta > 0$ that causes $y_j - c$ and $y_j^\theta - c^\theta$ to have different signs. When this is true, E_k^{obj} and E_k^{util} will also have different signs, resulting in the preference reversal observed in the –30 offset condition. There can be no such θ , however, because power functions increase monotonically for $\theta > 0$, meaning that utility cannot account for our data.

4.3. Serial order sequential sampling

The previous section illustrates that utility on its own cannot reverse the sign of the average difference between the uncertain and certain alternatives on –30 offset trials. Another way of accounting for different representations in the perceptual task and FGT is through serial-order effects, which discount a particular sample's contribution to the accumulated evidence based on when that sample was observed (Hogarth & Einhorn, 1992). These effects have been observed in experience-based preference tasks similar to the FGT (e.g., Hertwig, Barron, Weber, & Erev, 2004; Tsetos, Chater, & Usher, 2012).

Serial order effects can be formally incorporated into Eq. (2) by weighting the contribution of each sample j by an amount $a_j(k)$ between 0 and 1 whose value depends

on when j was observed. This results in the sequential sampling process

$$E_k^{\text{order}} = \sum_{j=1}^k a_j(k) \cdot (y_j^\theta - c^\theta). \quad (3)$$

Clearly, when $a_j(k) = 1$, sample j , $(y_j^\theta - c^\theta)$, contributes its full amount to the accumulated evidence, and when $a_j(k) = 0$, sample j does not contribute at all. When $a_j(k) < 1$, sample j contributes near its full amount when $a_j(k)$ is near one, and contributes nearly nothing when $a_j(k)$ is near zero. For example, a primacy effect can be accounted for by setting $a_1(k) = 1$ and $a_j(k) = 1/j$ and a recency effect by setting $a_k(k) = 1$ and $a_j(k) = 1/(k - j + 1)$ for $j < k$.

Again the only way to account for choices in the –30 uncertain offset condition is for $y_j - c$ and $a_j(k) \cdot (y_j^\theta - c^\theta)$ to have different signs. Following the previous section to rule out the utility parameter θ , the only way to reverse the sign is for the serial order weight $a_j(k)$ to be less than 0, $a_j(k) < 0$. However, if this were the case then the serial order updating model would allow for objective evidence in support of one alternative (e.g., uncertain alternative) to be represented as evidence for the other alternative (e.g., certain alternative). As this is not consonant with the usual definition of a serial order effect (Hogarth & Einhorn, 1992), we require $a_j(k) > 0$. By an argument analogous to the one presented for subjective expected utility, however, this means that incorporating serial order effects cannot account for choice data in the –30 uncertain offset condition.

4.4. Rank-dependent sequential sampling

4.4.1. Decision weights

A remaining hypothesis from behavioral decision theory is that people's attention to observed outcomes is not uniform, but varies with an outcome's likelihood of occurring and overall favorability (Tversky & Kahneman, 1992). This is typically captured using rank dependent utility theory and its decision weight construct (Luce, 2000; Quiggin, 1982; Tversky & Kahneman, 1992). According to rank dependent utility theory, the subjective value of the uncertain alternative is

$$v = \sum_{i=1}^n \pi_i x_i^\theta. \quad (4)$$

In this expression, $x_1 < \dots < x_n$ are the possible payoffs of the uncertain alternative ordered by their desirability and $\pi_1 < \dots < \pi_n$ are decision weights determining each payoff's contribution to the uncertain alternative's value. The decision weights π_i are a function of what Wakker (2010) calls “good-news probabilities” – the probabilities q_i of observing a payoff at least as large as x_i .

These probabilities are transformed into decision weights by a probability weighting function W . Following Prelec (1998), we use the probability weighting function with parameters γ and δ ,

$$W(q) = \exp[-\delta(-\ln(q))^\gamma]. \quad (5)$$

As illustrated in Fig. 5, γ controls the curvature of the weighting function, and δ controls the elevation of the




Fig. 5. Example probability weighting functions using the Prelec equation. (A) Illustration of how changes in the curvature parameter affect the shape of the function. (B) Illustration of how changes in the elevation parameter affect the shape of the function. Note the unity line (dotted line) denotes objective probability (i.e., no weighting).

function producing optimistic or pessimistic weights.¹ The decision weights π_i are determined by taking successive differences between the transformed good-news probabilities, so that $\pi_n = W(q_n)$ and $\pi_i = W(q_{i+1}) - W(q_i)$ for $i < n$. In the end, the decision weights reflect the marginal contribution of each possible outcome to the value of the alternative.²

4.4.2. Sequential sampling weights

A different formulation of the decision weights is needed in a sequential sampling process because this process is a running sum of sampled outcomes, not a weighted average as in Eq. (4). As a result, we need to model people's sensitivity to each sample as it is experienced. This sensitivity to each sample is approximately equal to the derivative of the probability weighting function,

$$\omega_i^* \approx \frac{W(q_{i+1}) - W(q_i)}{q_{i+1} - q_i} \quad (6)$$

(for a derivation of the sensitivity parameters ω_i^* see the Appendix). We show several probability weighting functions and their corresponding derivatives in Fig. 6. They illustrate how probability weighting functions color observed evidence. When $\omega_i^* > 1$, the decision maker shows greater sensitivity to the corresponding sample, giving it more weight than it objectively should be given. When $\omega_i^* < 1$, the decision maker is less sensitive to the sample. Finally, when $\omega_i^* = 1$, the decision maker treats the sample equivalent to its objective weight in the accumulated sum. Using these sensitivity weights, the evidence accumulation process is

$$E_k^{\text{weight}} = \sum_{j=1}^k \omega_j (y_j^\theta - c^\theta), \quad (7)$$

where ω_j is the ω_i^* associated with sample y_j . One interesting aspect of this rank-dependent sequential sampling model is that due to the weak law of large numbers, as sample size grows (e.g., due to larger choice thresholds), predicted choices from the model will converge with predicted choices from rank-dependent utility theory (see the Appendix for a formal proof). Thus, the rank-dependent sequential sampling model we developed here represents a dynamic and stochastic generalization of rank-dependent utility theory.

4.4.3. Optimistic sample weights

We assessed the ability of rank-dependent sequential sampling to account for differences between the FGT and perceptual tasks via simulation. We used simulation, rather than fitting the model, because the likelihood has not yet been derived for the rank-dependent sequential sampling model. Our simulation consisted of two stages. In the first stage, we determined the decision threshold A and perceptual measurement error τ , to be defined shortly, that provided the best fit to participants' behavior for each uncertain offset condition ($-30, 0, +30$) when the uncertain standard was 50 dots. A detailed description of the simulation-based fitting procedure is contained in the Appendix.

In the second stage, we simulated the sequential sampling rank dependent model using the values of A and τ from the perceptual task as estimates for the same parameters in the gambling task. Respectively, these estimates were 755 dots and 80 dots. This assumption seems reasonable given the fact that the analyses with the drift diffusion model did not yield significantly different thresholds for participants in the perceptual and gambling tasks. Moreover, linking the tasks in this way reduced the number of free parameters in our simulation while still allowing us to compare the FGT and perceptual analog.

¹ The curvature parameter is often described as controlling the overweighing and underweighing of rare events. However, in a rank dependent model the distortion is applied to the probability ranks so that large and small ranks are overweighed where the function is steepest (Wakker, 2010).

² The move to rank dependent calculation of decision weights over a direct transformation of probabilities (e.g., Kahneman & Tversky, 1979) address what appears to be implausible predictions of violation of stochastic dominance (Diecidue & Wakker, 2001).




Fig. 6. Probability weighting functions (A) and their corresponding derivatives (B). The derivatives of the functions roughly approximate the outcome weights used in the rank-dependent sequential sampling process.




Fig. 7. Simulation results from the rank-dependent sequential sampling process. (A) Values of δ and γ that satisfied two conditions, colored by the variability of their RTs across the three offset levels. The first condition was that the uncertain alternative was chosen on at least half of trials for all offset levels. The second was that the number of times the uncertain alternative was chosen must increase with increasing offset. (B) Prelec functions corresponding to the δ and γ values in A. (C) The derivatives of the Prelec functions in B. The lines in B and C were colored by the variability of their RTs across the three offset levels.

Each decision was simulated by sequentially sampling outcomes from the uncertain alternative, perturbing these outcomes with normally-distributed noise,³ inserting the noisy values into Eq. (7) and making a decision when E_k^{weight} reached the corresponding threshold. The noise simulated perceptual measurement error, and its standard deviation is the free parameter τ estimated in the perceptual task. For simplicity, we set the utility parameter θ in Eq. (7) to one. To investigate the sequential sampling weights that best recreate the gambling data we simulated 1000 decisions for each γ and δ in a grid of values extending from 0 to 2.⁴ We ran this simulation for the experimental conditions in the FGT that corresponded to the conditions that informed the perceptual model (i.e., the offset conditions of –30, 0 and +30 dots, when the standard deviation was 50 dots).

Fig. 7 summarizes the results of this simulation. Panel A shows which values of γ and δ resulted in choice proportions and RTs satisfying the following three conditions.

³ We also took the floor of the noisy value, so the noisy value still corresponded to a number of dots.

⁴ Specifically, we simulated data for every (γ, δ) such that $\gamma = k_1/20$ and $\delta = k_2/20$, where $1 \leq k_1, k_2 \leq 40$. This means that $1/20 \leq \gamma, \delta \leq 2$.

First, choice proportions in the –30 offset condition must be above 0.5. Second, choice proportions must monotonically increase with increasing uncertain offset. Finally, we calculated the mean-squared deviation between the observed and predicted RTs over all decisions, show only those values whose deviation was less than 0.5. The shading in Fig. 7 shows the mean-squared deviation for the points satisfying these conditions.

Inspection of the plot reveals that γ values near 0.7 and δ values near 0.3 do a good job of recreating our data. Panel B plots probability weights ($W(q_i)$) for the values of γ and δ shown in Panel A, and Panel C plots the sequential sampling weights (ω_i^*). These results show that probability weighting can indeed qualitatively account for our data. Moreover, with $\gamma, \delta = 1$ in the perceptual task, the best fitting values for the FGT suggest that participants weighted observed outcomes more optimistically than participants in the perceptual task. In particular, participants in the FGT tended to emphasize larger gains and deemphasize smaller gains in determining the value of the uncertain alternative.

Apparently, simply framing the task as a gamble is sufficient to change how decision makers allocate attentional weight to sampled outcomes. Non-linear decision weights

have sometimes been explained in terms of a perceptual-like distortion (Kahneman & Tversky, 1984). Indeed, choices in the perceptual condition – specifically the bias to choose the uncertain alternative when the two alternatives had equal means – is consistent with some over-weighting of large magnitudes. However, the comparison between the FGT and its perceptual analog suggest the rank dependent weights go beyond any simple constant level of perceptual distortion. Similar findings of decision makers overweighting extreme payoffs have been reported in other studies of decisions made from experience (Ludvig, Madan, & Spetch, 2013; Tsetsos et al., 2012) and we discuss this finding in the context of those studies. Before doing so, we discuss the implications of our results for computational models of preferential choice.

4.5. Implications for computational models of preferential choice

Besides revealing some of the basic properties of rapid decisions from experience, our results are also informative for computational models of decision making (Busemeyer & Townsend, 1993; Usher & McClelland, 2004). As discussed before, these models often assume a sequential sampling process underlies preferential choice. This sequential sampling assumption implies that to make a choice, decision makers mentally sample or simulate possible outcomes from the available alternatives. These outcomes could be memory traces of past outcomes or, in the case of gambles, they could be newly sampled outcomes as one's attention is drawn to consider different possible outcomes (Busemeyer, 1982).

One assumption in these models is that the contribution of the sampled outcome to the accumulated evidence does not depend on the magnitude of the sampled outcome. Our results suggest otherwise. The rank-dependent sample weights imply the contribution of each outcome to the accumulated evidence at each time point is dependent partly on the relative standing of the sampled outcome in the larger set of possible outcomes. Having said that, rank-dependent sample weights does not discount the role of other factors, such as serial dependency (e.g., Tsetsos et al., 2012) and competing interactions with other alternatives (Roe, Busemeyer, & Townsend, 2001), in shaping preferential choice. In fact, the FGT and similar tasks (Tsetsos et al., 2012) make the systematic investigation of these and other factors possible, through their ability to explicitly control the flow of samples of reward information.

4.6. Decisions from experience

Our study falls into the general category of experience-based decisions where people must gather information from noisy samples of outcomes (e.g., Busemeyer, 1985; Glaser, Trommershäuser, Mamassian, & Maloney, 2012; Hertwig et al., 2004; Ungemach, Chater, & Stewart, 2009). The FGT and our results contribute to our general understanding of experience-based decisions in several different areas. Methodologically, the FGT expands the decisions-from-experience paradigm in at least three different areas.

First, it expands the paradigm to rapid presentations of payoff information. Typically the paradigm has focused on slow trial-by-trial presentation of payoff information (e.g., Barron & Erev, 2003; Hertwig et al., 2004). In a notable extension of the paradigm, Tsetsos et al. (2012) increased the presentation rate to 2–4 Hz; here we have pushed it even further to 20 Hz (i.e., a sample every 50 ms). This means that for the average decision, which participants took about 1.37 s to make, they were shown about 27 samples from the uncertain alternative. Compare that with the slower experience-based decision paradigms where participants often base their decisions on 7–10 observations from a given alternative with each observation being displayed for a second or two (Hertwig et al., 2004). Investigations of how people make these rapid decisions from experience is important. One reason is that there is a vast array of preferential decisions made on the basis of rapidly communicated statistical information of which our day-trader example in the introduction is just one example. Our work shows that these decisions are well described as a sequential sampling process, a process well known in statistics (Wald & Wolfowitz, 1949) and used to make other more perceptual (Gold & Shadlen, 2007; Link & Heath, 1975) and memory-based decisions (Ratcliff, 1978).

A second methodological contribution is the extension of experience-based decisions to continuous distributions of payoffs. Most current studies of experience-based decisions focus on binary gambles whose payoffs follow a discrete probability distribution (Fox & Hadar, 2006; Gonzalez & Dutt, 2011; Hau, Pleskac, Hertwig, 2010; Hau, Pleskac, Kiefer, & Hertwig, 2008; Hertwig et al., 2004; Hills & Hertwig, 2010; Ludvig et al., 2013; Rakow, Demes, & Newell, 2008; Rakow & Newell, 2010; Ungemach et al., 2009; for studies that have used multiple outcome gambles see Busemeyer (1985), Ert & Erev (2007), Thaler, Tversky, Kahneman, & Schwartz (1997) and Barron & Erev (2003)). Discrete binary gambles are a useful simplification, but they have constrained the questions being asked. For example, to date most of the work on experience-based choice has centered around the height of the probability weighting function at specific probability values (Hertwig & Erev, 2009). Moving to a continuous distribution forced us to consider the implications of experiencing a range of payoffs and the contextual effects of those payoffs on each sample of experience. Here we found these contextual effects were well captured with a weighting function adapted from rank-dependent utility theories (Luce, 2000; Quiggin, 1982; Tversky & Kahneman, 1992; Wakker, 2010) within a sequential sampling framework.

A third methodological contribution is that the perceptual analog to the FGT provides a different perspective for understanding experience-based preferential choice. Recently, many studies have sought to understand decisions from experience by comparing them with so-called decisions from description, in which respondents learn about the outcomes and probabilities of gambles through convenient descriptions. The main finding that comes from this comparison is that in experience-based choices people choose as if the impact of objectively rare outcomes has been attenuated in comparison to the weight given to the same outcomes in decisions from description producing a

so-called description-experience gap (Hau, Pleskac, Kiefer et al., 2008).⁵ This description control, however, has proven troublesome in trying to understand the description-experience gap and more generally how people make decisions from experience. One reason is that the comparison between decisions from description and decisions from experience confounds sampling error, which alone can produce underweighting of rare events (Fox & Hadar, 2006; Rakow et al., 2008; Ungemach et al., 2009). That is, during decisions from experience one factor that impacts the internal representation of the gambles is the random variability in the sampled payoffs, which is not present when choices are made from descriptions of the gambles. Our comparison between preferential and perceptual choice using the FGT and its perceptual analog controls for sampling error, since both tasks are experience-based and use the same information, but framed differently. Note also that sampling error should also be reduced via the larger sample sizes that are taken to make each decision in the FGT. The perceptual task also controls for other confounds that the comparison to decisions from description does not. One of these is the perceptual error that naturally arises when outcomes/payoffs are conveyed via experience as compared to description (Shafir, Reich, Tsur, Erev, & Lotem, 2008). Other confounds include how respondents search or sample the alternatives for information, the memory requirements, and even possibly the decision process itself (Hertwig & Erev, 2009).

Together these methodological extensions helped reveal that while people use similar cognitive machinery to make rapid experience-based perceptual and preferential choices, the underlying processes also systematically differ. Specifically, the valuation process used in these preferential decisions appears to change the subjective representation of the statistical information with more attentional weight being placed on larger and more extreme potential payoffs. This can lead to risk seeking preferences in the gain domain when the risky option has extreme potential payoffs. Our result echoes similar results found in other studies with experience-based decision making (Ludvig et al., 2013; Tsetsos et al., 2012).⁶ The next question is why do people overweight these extreme events. Ludvig et al. (2013) suggests that this effect is similar to perceptual context effects where, for instance, colors appear differently depending on the other colors around them (Lotto & Purves, 2000) or lines look longer or shorter depending on their arrowheads (Müller-Lyer, 1889). In the case of gambles, payoffs can seem more or less extreme relative to other

possible payoffs. In our case, the rank-dependent nature of the weights suggests this arises because of the range of payoffs people experience during a particular trial with the more extreme payoffs grabbing more attention. Our perceptual control, however, helps rule out a purely perceptual explanation to this effect. Instead, as mentioned earlier, it appears that framing the task as a gamble changes the attentional weight allocated to possible outcomes and the fact that this salience arises at rapid presentation rates would seem to speak against more deliberative processes like learning (Niv, Edlund, Dayan, & O'Doherty, 2012). Perhaps the change in attention is the result of a motivational (Lopes, 1987; Weber, 1994) or arousal (Pham, 2007) component that arises in making preferential decisions. Either way our results support a growing appreciation of the role that attention plays in forming a preference during economic decisions (Krajbich, Armel, Rangel et al., 2010; Krajbich & Rangel, 2011).

5. Conclusion

In conclusion, in this paper we investigated rapid experience-based decision making using a novel gambling task, the FGT. This task and its perceptual analog showed that people use a similar sequential sampling process to make these two decisions. However, there are critical differences that lead to a dissociation between preferential and perceptual choice. We show that the valuation process used during preferential choices changes the representation of the statistical information that is accumulated. In particular, the results are consistent with a sequential sampling process using rank dependent weights where decision makers optimistically weighted higher payoffs during deliberation.

Acknowledgements

We would like to thank Mitchell Uitvlugt for his assistance in preparing the manuscript. We would also like to thank the members of the Laboratory for Cognitive and Decision Science at Michigan State University for their insightful comments on earlier drafts of this paper.

This work was supported in part by NIH Grant R03DA033455 awarded to TJP and TL.

Appendix A. Derivation of sensitivity parameters in rank-dependent sequential sampling process

A.1. Definition of the sensitivity parameter

The sensitivity weights in the rank-dependent sequential sampling process were formed in a manner consistent with rank dependent utility theory (Wakker, 2010, Appendix 6.8). This supplement describes how they were derived and how the model was estimated. Recall that the value a decision maker assigns to an uncertain alternative under rank-dependent utility theory is

$$v = \sum_{i=1}^n \pi_i x_i^{\theta}, \quad (\text{A.1})$$

⁵ In some cases, the weight given to rare outcomes is attenuated relative to the weight they deserve according to their objective probabilities (see for example Problems 5 and 6 in Hertwig et al., 2004).

⁶ This result of greater weight to more extreme payoffs does seem to contrast with the typical finding of underweighting of rare events often reported in decisions from experience (Hertwig & Erev, 2009). However, as described earlier, the underweighting of rare events is often defined relative to the choices from decisions from description. Our study, as we have discussed, uses a perceptual control. Even so, it may be the case that underweighting of rare events is confined to the experience-based paradigms that use slow trial-by-trial presentation of payoff information. As we have described in the discussion, there are a number of important differences like the rapid presentation of larger sample sizes on each trial that could lead to this difference. This is certainly an avenue for future research.

where θ is the exponential parameter in the utility function (see Section 3), $x_1 < \dots < x_n$ are the possible outcomes of the uncertain alternative ordered by desirability and each π_i is the decision weight assigned to x_i . The π_i are determined by the good-news probabilities q_i and the [Prelec \(1998\)](#) function

$$W(q) = \exp[-\delta(-\ln(q))^\gamma] \quad (\text{A.2})$$

through the equation

$$\pi_i = W(q_{i+1}) - W(q_i). \quad (\text{A.3})$$

Note that $p_i = q_{i+1} - q_i$. Since the Prelec function is differentiable,⁷

$$\pi_i = W(q_{i+1}) - W(q_i) = \int_{q_i}^{q_{i+1}} W'(\tilde{q}) d\tilde{q} \quad (\text{A.5})$$

by the Fundamental Theorem of Calculus (see [Wakker, 2010, p. 199, for an in-depth discussion](#)). Multiplying by $p_i/(q_{i+1} - q_i) = 1$, gives

$$\pi_i = p_i \times \frac{\int_{q_i}^{q_{i+1}} W'(\tilde{q}) d\tilde{q}}{q_{i+1} - q_i}. \quad (\text{A.6})$$

The ratio on the right-hand side of the product in Eq. (A.6) is the average slope over the interval $[q_i, q_{i+1}]$ and is approximately equal to the derivative of the weighting function when $p_i = q_i - q_{i+1}$ is small. This factor measures the sensitivity of the overall rank dependent value to each outcome. We set the outcome weights in the sequential sampling process to this sensitivity term,

$$\omega_i^* = \frac{\int_{q_i}^{q_{i+1}} W'(\tilde{q}) d\tilde{q}}{q_{i+1} - q_i}. \quad (\text{A.7})$$

A.2. Agreement with rank-dependent utility

Describing π_i in this way allows ω_i^* to be incorporated into a sequential sampling process as

$$E_k^{\text{weight}} = \sum_{j=1}^k \omega_j (y_j^\theta - c^\theta), \quad (\text{A.8})$$

where ω_j is the ω_j^* associated with outcome y_j . Note that Eq. (A.8) is the same as Eq. (7) except that, Eq. (7) contains a perceptual error term and Eq. (A.8) does not. In this section, we show that when the decision criterion is sufficiently large, decisions made by the sequential sampling process defined Eq. (A.8) agree with the decisions made by rank-dependent utility theory.⁸ This is sufficient to show that the sequential sampling process defined in Eq. (7) is a process-model generalization of rank-dependent utility theory.

As noted earlier, this is a consequence of the Weak Law of Large Numbers (WLLN; [Casella & Berger, 2002](#)). The

WLLN formally states that, in a large sample, the proportion of times a outcome x_i occurring with probability $p_i = q_{i+1} - q_i$ will be observed is approximately its probability $q_{i+1} - q_i$. For any sample of size k ,

$$E_k^{\text{weight}} = \sum_{j=1}^k \omega_j (y_j^\theta - c^\theta) = \sum_{i=1}^n m_{i,k} \omega_i^* (x_i^\theta - c^\theta), \quad (\text{A.9})$$

where $m_{i,k}$ is the number of times outcome i is observed the sample. When k is sufficiently large $m_{i,k} \omega_i^* \approx k(q_{i+1} - q_i) \pi_i$ by the WLLN. Substituting $k(q_{i+1} - q_i) \pi_i$ into Eq. (A.9), yields

$$E_k^{\text{weight}} \approx k \sum_{i=1}^n \pi_i (x_i^\theta - c^\theta) = k v. \quad (\text{A.10})$$

When the decision threshold is large, a large number of samples will be observed before a decision is made. Thus, for large thresholds, the average evidence accumulated for the uncertain alternative, E_k^{weight}/k , is approximately the value v assigned to the uncertain alternative by rank-dependent utility theory. Thus, the decisions made by the sequential sampling model and rank dependent utility will agree.

Appendix B. Fitting the simulation free parameters

To fit the decision threshold A and perceptual measurement error τ , we estimated a decision threshold and perceptual measurement error for each of the three uncertain offset conditions ($-30, 0, +30$) when the uncertain standard deviation was 50 dots using participants' behavior in the perceptual task. We then averaged these estimates across all uncertain offset conditions to obtain the estimates used to simulate decisions using rank-dependent weights.

Because the best-fitting decision threshold depends on the amount of perceptual measurement error, we fit each condition's decision threshold and perceptual measurement error in two steps. First, we fit a decision threshold for each measurement error in the set $\{1, 10, \dots, 100\}$. This was necessary because the best-fitting decision threshold depends on the level of perceptual measurement error. Then, we visually compared the choice and RT predictions for each decision threshold/perceptual measurement error pair, and selected the one that best captured participant's behavior in the perceptual task. The resulting error estimates were identical across the three tasks (see [Table B.1](#)).

To fit the decision threshold given a level of perceptual measurement error, we first computed the average number of samples viewed on a trial, M , which is equal to the mean RT in seconds divided by the rate at which samples

Table B.1

Best-fitting τ and A values as a function of the uncertain offset μ .

Uncertain offset (μ)	Measurement error ($\hat{\tau}_c$)	Threshold (A_c)
−30	80 dots	1030 dots
0	80	227
+30	80	1012
Mean	80	755

⁷ The derivative of the [Prelec \(1998\)](#) weighting function is

$$W'(q) = \frac{\gamma \delta}{q} \exp[-\delta(-\ln(q))^\gamma] [-\ln(q)]^{\gamma-1}. \quad (\text{A.4})$$

⁸ Formally, the limit as the threshold increases, the probability of disagreement converges to zero.

are drawn. We then simulated $N = 1000$ sequences of M independent dot differences $y_{i1}, \dots, y_{iM}$ and measurement errors $\epsilon_{i1}, \dots, \epsilon_{iM}$ for each condition, and estimated the condition's threshold as the average total evidence collected in the average number of samples collected in a trial,

$$\hat{A}_c = \frac{1}{N} \sum_{i=1}^N \left| \sum_{j=1}^M y_{ij} + \epsilon_{ij} \right|. \quad (\text{B.1})$$

We then estimated the threshold parameter as the average of the A_c . The values of A_c fit for each condition are shown in Table B.1.

Examination of Table B.1 shows that the thresholds fit for the ± 30 offset conditions are approximately equal to one another, but not to the threshold fit for the 0 offset condition. This is a consequence of all three conditions having the same mean RT despite their difference in mean perceptual evidence. Because they have the same mean RT the mean number of observed samples M will also be the same. As a result, the threshold estimated in the 0 offset condition must be smaller than the thresholds in the ± 30 condition since the absolute value of the evidence collected from each sample is smaller in the 0 condition (i.e., on average the value of the additional evidence will be 0). To model the data in the gambling trials, we used a single threshold, by taking the mean of threshold estimates for the three offset conditions for the perceptual trials (i.e., 755 dots). In addition, given the large difference between the 0 and ± 30 conditions, we also ran a simulation using different thresholds for each condition to make sure that our results were not simply an artifact of averaging the thresholds. Though not presented here, we obtained qualitatively similar results from both simulations.

References

- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215–233.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22, 23–36.
- Busemeyer, J. R. (1982). Choice behavior in a sequential decision-making task. *Organizational Behavior & Human Performance*, 29, 175–207.
- Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 11, 538–564.
- Busemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. Thousand Oaks, CA: SAGE Publications.
- Busemeyer, J. R., Jessup, R. K., Johnson, J. G., & Townsend, J. T. (2006). Building bridges between neural models and complex decision making behaviour. *Neural Networks*, 19, 1047–1058.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100, 432–459.
- Casella, G., & Berger, R. L. (2002). *Statistical inference* (2nd ed.). Pacific Grove, CA: Duxbury.
- Diecidue, E., & Wakker, P. P. (2001). On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23, 281–298.
- Diederich, A., & Busemeyer, J. R. (1999). Conflict and the stochastic-dominance principle of decision making. *Psychological Science*, 10, 353–359.
- Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making*, 20, 305–322.
- Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = Sampling error plus prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making Journal*, 1, 159–161.
- Glaser, C., Trommershäuser, J., Mamassian, P., & Maloney, L. T. (2012). Comparison of the distortion of probability information in decision under risk and in an equivalent visual task. *Psychological Science*, 23, 419–426.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574.
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118, 523.
- Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making*, 23, 48–68.
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21, 493–518.
- Heathcote, A., Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, 9, 394–401.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539.
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Science*, 13, 517–523.
- Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, 21, 1787–1792.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating – The belief-adjustment model. *Cognitive Psychology*, 24, 1–55.
- Hothorn, T., Bretz, F., & Westfall, P. (2008). Simultaneous inference in general parametric models. *Biometrical Journal*, 50, 346–363.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–292.
- Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, 39, 341–350.
- Krajibich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13, 1292–1298.
- Krajibich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, 3.
- Krajibich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 13852–13857.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. New York, NY: Academic Press.
- Link, S. W., & Heath, R. A. (1975). Sequential theory of psychological discrimination. *Psychometrika*, 40, 77–105.
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6, 312–319.
- Lopes, L. L. (1987). Between hope and fear – The psychology of risk. *Advances in Experimental Social Psychology*, 20, 255–295.
- Lotto, R. B., & Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of the National Academy of Sciences*, 97, 12834–12839.
- Luce, R. D. (2000). *Utility of gains and losses: Measurement-theoretical and experimental approaches*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2013). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*. <http://dx.doi.org/10.1002/bdm.1792>.
- Milosavljevic, M., Malmud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making*, 5, 437–449.
- Müller-Lyer, F. C. (1889). Optische urteilstäuschungen. *Archiv für Anatomie und Physiologie, Physiologische Abteilung*, 2, 263–270.
- Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *The Journal of Neuroscience*, 32, 551–562.
- Palmer, J., Huk, A., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5, 376–404.
- Pham, M. T. (2007). Emotion and rationality: A critical review and interpretation of empirical evidence. *Review of General Psychology*, 11, 155–178.
- Prelec, D. (1998). The probability weighting function. *Econometrica*, 66, 497–527.
- Quiggin, J. (1982). A theory of anticipated risk. *Journal of Economic Behavior & Organization*, 3, 323–343.

- Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior & Human Decision Processes*, 106, 168–179.
- Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, 23, 1–14.
- Ratcliff, R. (1978). Theory of memory retrieval. *Psychological Review*, 85, 59–108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20, 873–922.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108, 370–392.
- Savage, L. J. (1954). *The foundations of statistics*. New York, NY: John Wiley & Sons.
- Shadlen, M. N., Kiani, R., Hanks, T., & Churchland, A. K. (2008). Neurobiology of decision making: An intentional framework. In C. Engel & W. Singer (Eds.), *Better than conscious? Decision making, the human mind, and implications for institutions*. Cambridge, MA: MIT Press.
- Shafir, S., Reich, T., Tsur, E., Erev, I., & Lotem, A. (2008). Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature*, 453, 917–920.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25, 251–260.
- Summerfield, C., & Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: Neural and computational mechanisms. *Frontiers in Neuroscience*, 6.
- Symmonds, M., & Dolan, R. J. (2012). The neurobiology of preference. In R. J. Dolan & T. Sharot (Eds.), *Neuroscience of preference and choice: Cognitive and neural mechanisms* (pp. 4–23). Waltham, MA: Elsevier.
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *Quarterly Journal of Economics*, 112, 647–661.
- Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109, 9659–9664.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, 20, 473–479.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111, 757–769.
- von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.
- Wagenmaker, E.-J., Krypotos, A. M., Criss, A. H., & Iverson, G. (2012). On the interpretation of removable interactions: A survey of the field 33 years after Loftus. *Memory & Cognition*, 40, 145–160.
- Wakker, P. P. (2010). *Prospect theory for risk and ambiguity*. Cambridge, UK: Cambridge University Press.
- Wald, A., & Wolfowitz, J. (1949). Bayes solutions of sequential decision problems. *Proceedings of the National Academy of Sciences of the United States of America*, 35, 99–102.
- Weber, E. U. (1994). From subjective probabilities to decision weights: The effect of asymmetric loss functions on the evaluation of uncertain outcomes and events. *Psychological Bulletin*, 115, 228–242.
- Weber, E. U., Shafir, S., & Blais, A. R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, 111, 430–445.